

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

AF
JFW

Applicant: Gregory Marlan et al.

Title: DETECTION AND CONTROL OF RESOURCE CONGESTION BY A NUMBER OF PROCESSORS

Docket No.: 499.750US1
Filed: July 31, 2003
Examiner: Ryan A. Dare

Serial No.: 10/631,988
Due Date: August 17, 2007
Group Art Unit: 2186

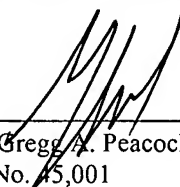
MS Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

We are transmitting herewith the following attached items (as indicated with an "X"):

- ☒ Return postcard.
- ☒ Response to Notification of Non-Compliant Appeal Brief (Supplemental Appeal Brief Under 37 C.F.R. 41.37) (31 pgs.).

If not provided for in a separate paper filed herewith, Please consider this a **PETITION FOR EXTENSION OF TIME** for sufficient number of months to enter these papers and please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
Customer Number 21186

By: 
Atty: Gregg A. Peacock
Reg. No. 45,001

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: MS Amendment, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 9 day of August, 2007.

Lisa Rosarske
Name

Lisa Rosarske
Signature

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
(GENERAL)



SUPPLEMENTAL APPEAL BRIEF UNDER 37 C.F.R. § 41.37

TABLE OF CONTENTS

	<u>Page</u>
<u>1. REAL PARTY IN INTEREST</u>	2
<u>2. RELATED APPEALS AND INTERFERENCES</u>	3
<u>3. STATUS OF THE CLAIMS</u>	4
<u>4. STATUS OF AMENDMENTS</u>	5
<u>5. SUMMARY OF CLAIMED SUBJECT MATTER</u>	6
<u>6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL</u>	10
<u>7. ARGUMENT</u>	11
<u>8. SUMMARY</u>	21
<u>CLAIMS APPENDIX</u>	22
<u>EVIDENCE APPENDIX</u>	30
<u>RELATED PROCEEDINGS APPENDIX</u>	31



PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gregory Marlan et al.

Examiner: Ryan A. Dare

Serial No.: 10/631,988

Group Art Unit: 2186

Filed: July 31, 2003

Docket: 499.750US1

For: DETECTION AND CONTROL OF RESOURCE CONGESTION BY A NUMBER
OF PROCESSORS

SUPPLEMENTAL APPEAL BRIEF UNDER 37 CFR § 41.37

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

The Supplemental Appeal Brief is in response to the Notification of Non-Compliant Appeal Brief dated July 17, 2007, and the Non-Final Office Action dated March 23, 2007, and is presented in support of the Notice of Appeal to the Board of Patent Appeals and Interferences, filed on July 18, 2006, from the Final Rejection of claims 1-40 of the above-identified application, as set forth in the Final Office Action mailed on April 18, 2006.

The Appellants respectfully request consideration and reversal of the Examiner's rejections of pending claims.

(
1. REAL PARTY IN INTEREST

The real party in interest of the above-captioned patent application is the assignee,
SILICON GRAPHICS, INC.

2. RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to Appellant that will have a bearing on the Board's decision in the present appeal.

3. STATUS OF THE CLAIMS

The present application was filed on July 31, 2003 with 40 claims. In response to the Office Action mailed October 28, 2005, an Amendment was filed to amend claim 27 and 34-40. Claims 1-40 stand rejected, remain pending, and are the subject of the present Appeal.

4. STATUS OF AMENDMENTS

No amendments have been made subsequent to the Final Office Action dated April 18, 2006.

5. SUMMARY OF CLAIMED SUBJECT MATTER

Some aspects of the present inventive subject matter include, but are not limited to, methods, systems and apparatus for detection and control of resource congestion by a number of processors. (See Figure 1).

In claim 1, an apparatus comprises a load/store unit that includes a retry logic that is to retry access to a resource after receipt of a negative acknowledgement for an attempt to access the resource by the load/store unit. (See the load/store functional unit 218 and the retry logic 286 in Figure 2 and page 15, line 27 – page 16, line 6). The apparatus also includes a congestion detection logic to output a signal that indicates that the resource is congested based on receipt of a consecutive number of negative acknowledgments in response to access requests to the resource. (See congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3).

In claim 5, a processor comprises a functional unit to attempt to access data from memory coupled to the processor based on an access request. (See the load/store functional unit 218 in Figure 2 and page 15, line 27 – page 16, line 6). The functional unit is to retry attempts to access of the data based on other access requests after receipt of a negative acknowledgement in response to the attempt to access the data. (See the load/store functional unit 218 in Figure 2 and page 15, line 27 – page 16, line 6). The processor also comprises a congestion detection logic to detect congestion of access of the data based on receipt of a consecutive number of negative acknowledgments that exceed a threshold prior to access of the data. (See the congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3). The processor comprises a congestion control logic to disable the functional unit from the attempts to access the data for a time period after congestion is detected. (See the congestion control logic 284 in Figure 2 and page 16, lines 7-13).

In claim 8, a processor comprises a functional unit to attempt to access a cache line in a cache memory coupled to the processor based on an access request. (See the processors 104 and the caches 102 in Figure 1 and page 13, lines 4-26). The functional unit is to retry attempts to access the cache line based on additional access requests after receipt of a negative acknowledgement in response to the attempt to access the data. (See the load/store functional unit 218 in Figure 2 and page 15, line 27 – page 16, line 6). The processor also comprises a

congestion detection logic to detect congestion of access of the cache line based on an average number of negative acknowledgments received that exceed a threshold prior to access of the data. (See the congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3). The processor comprises a congestion control logic to disable the functional unit from attempts to access the cache line for a time period after congestion is detected. (See the congestion control logic 284 in Figure 2 and page 16, lines 7-13).

In claim 12, a system comprises a cache memory to store data. (See the caches 102 in Figure 1 and page 12, lines 20-26). The system also comprises a first processor to attempt to access the data from the cache memory based on access requests. (See the processors 104 in Figure 1 and page 13, lines 4-26). The first processor includes a congestion detection logic to detect congestion of access to the data based on receipt of a consecutive number of negative acknowledgements in response to the access requests. (See the congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3).

In claim 17, a system comprises a resource. (See the caches 102 in Figure 1 and page 12, lines 20-26). The system also comprises a first processor having a load/store functional unit. (See the processor 104 and the load/store functional unit 218 in Figure 2 and page 15, lines 24-27). The load/store functional unit is to attempt to access the resource based on access requests. (See the load/store functional unit 218 in Figure 2 and page 15, line 27 – page 16, line 6). The first processor includes a congestion detection logic to detect congestion of access of the resource based on a consecutive number of negative acknowledgements received in response to the access requests prior to receipt of a positive acknowledgment in response to one of the access requests within a first time period. (See the congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3).

In claim 22, a system comprises a cache memory to include a number of cache lines for storage of data. (See the caches 102 in Figure 1 and page 12, lines 20-26). The system also comprises at least two processors, wherein a first processor of the at least two processors is to attempt to access the data in one of the number of cache lines based on access requests. (See the processors 104 in Figure 1 and page 13, lines 4-26). The first processor includes a congestion detection logic to detect congestion of access of a first cache line of the number of cache lines based on a ratio of a number of negative acknowledgments to a number of positive

acknowledgments received in response to the access requests. (See the congestion detection logic 282 in Figures 2 and 3 and page 16, line 24 – page 3).

In claim 27, a method comprises transmitting access requests, by a first processor, to access data in a memory. (See the block 702 in Figure 7 and page 20, lines 7-14). The method also comprises receiving a positive acknowledgement or a negative acknowledgment from a second processor that is associated with the memory based on one of the number of access requests. (See the block 704 in Figure 7 and page 20, lines 15-22). The method comprises detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment. (See the block 710 in Figure 7 and page 21, lines 13-21; also see the block 912 in Figure 9 and page 24, line 25 – page 25, line 7).

In claim 31, a method comprises accessing, by at least one processor, a resource based on an access request. (See the block 902 in Figure 9 and page 23, lines 3-6). The method also comprises receiving a positive acknowledgement if the resource is accessible. (See the block 906 in Figure 9 and page 23, lines 7-15). The method comprises receiving a negative acknowledgement if the resource is not accessible. (See the block 906 in Figure 9 and page 23, lines 7-15). The method comprises retrying accessing, by the at least one processor, of the resource based on a number of access requests. (See the block 910 in Figure 9 and page 23, line 20 – page 24, line 24). The method also comprises detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time period, prior to receiving a positive acknowledgments. (See the block 910 in Figure 9 and page 23, line 20 – page 24, line 24).

In claim 34, a machine-readable medium provides instructions, which when executed by a machine, cause said machine to perform operations. (See page 11, lines 1-10). The operations comprise transmitting access requests, by a first processor, to access data in a memory. (See the block 702 in Figure 7 and page 20, lines 7-14). The operations also comprise receiving a positive acknowledgement or a negative acknowledgment from a second processor that is associated with the memory based on one of the number of access requests. (See the block 704 in Figure 7 and page 20, lines 15-22). The operations also comprise detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative

acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment. (See the block 710 in Figure 7 and page 21, lines 13-21).

In claim 38, a machine-readable medium provides instructions, which when executed by a machine, cause said machine to perform operations. (See page 11, lines 1-10). The operations comprise accessing, by at least one processor, a resource based on an access request. (See the block 902 in Figure 9 and page 23, lines 3-6). The operations comprise receiving a positive acknowledgement if the resource is accessible. (See the block 906 in Figure 9 and page 23, lines 7-15). The operations also comprise receiving a negative acknowledgement if the resource is not accessible. (See the block 906 in Figure 9 and page 23, lines 7-15). The operations comprise retrying accessing, by the at least one processor, of the resource based on a number of access requests. (See the block 910 in Figure 9 and page 23, line 20 – page 24, line 24). The operations comprise detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time period, prior to receiving a positive acknowledgments. (See the block 910 in Figure 9 and page 23, line 20 – page 24, line 24).

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-40 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Hughes et al. (U.S. Patent No. 6,427,193) (hereinafter “Hughes”) in view of Swami (U.S. Patent Application No. 2004/0165538) (hereinafter “Swami”).

7. ARGUMENT

A) The Applicable Law under 35 U.S.C. §103

The Examiner has the burden under 35 U.S.C. § 103 to establish a *prima facie* case of obviousness.¹ First and foremost, to establish *prima facie* obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art.² In combining prior art references to construct a *prima facie* case, the Examiner must show some objective teaching in the prior art or some knowledge generally available to one of ordinary skill in the art that would lead an individual to combine the relevant teaching of the references.³ The M.P.E.P. contains explicit direction to the Examiner that agrees with the *In re Fine* court:

In order for the Examiner to establish a *prima facie* case of obviousness, three base criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure.⁴

An invention can be obvious even though the suggestion to combine prior art teachings is not found in a specific reference.⁵ However, while it is not necessary that the cited references or prior art specifically suggest making the combination, there must be some teaching somewhere which provides the suggestion or motivation to combine prior art teachings and applies that combination to solve the same or similar problem which the claimed invention addresses. One of ordinary skill in the art will be presumed to know of any such teaching.⁶ However, the level of skill is not that of the person who is an innovator but rather that of the person who follows the conventional wisdom in the art.⁷ The requirement of a suggestion or motivation to combine references in a *prima facie* case of obviousness is emphasized in the Federal Circuit opinion, *In*

¹ *In re Fine*, 837 F.2d 1071, 1074, 5 U.S.P.Q.2d (BNA) 1596, 1598 (Fed. Cir. 1988).

² *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974); M.P.E.P. § 2143.03.

³ *In re Fine* at 1598.

⁴ M.P.E.P. § 2142 (citing *In re Vaeck*, 947 F.2d 488, 20 U.S.P.Q.2d (BNA) 1438 (Fed. Cir. 1991)).

⁵ *In re Oetiker*, 977 F.2d 1443, 24 U.S.P.Q.2d (BNA) 1443 (Fed. Cir. 1992).

⁶ (See, e.g., *In re Nilssen*, 851 F.2d 1401, 1403, 7 U.S.P.Q.2d 1500, 1502 (Fed. Cir. 1988) and *In re Wood*, 599 F.2d 1032, 1037, 202 U.S.P.Q. 171, 174 (C.C.P.A. 1979)).

⁷ *Standard Oil Co. v. American Cyanamid Co.*, 774 F.2d 448, 474, 227 U.S.P.Q. 293, 298 (Fed. Cir. 1985).

re Sang Su Lee, 277 F.3d 1338; 61 U.S.P.Q.2D 1430 (Fed. Cir. 2002), which notes that the motivation must be supported by evidence in the record.

The test for obviousness under § 103 must take into consideration the invention as a whole; that is, one must consider the particular problem solved by the combination of elements that define the invention.⁸ References must be considered in their entirety, including parts that teach away from the claims.⁹ The fact that references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination.¹⁰

Recently, the Supreme Court reaffirmed the validity of the “teaching, suggestion, motivation” test in *KSR Int’l Co. v. Teleflex Inc.*, No. 04-1350 (U.S. Apr. 30, 2007) and guidance provided in a PTO Memo of May 3, 2007 recognizes this holding. In addition, the PTO Memo of May 3, 2007 indicated that “analysis supporting a rejection under 35 U.S.C. § 103(a) should be made explicit,” citing the Court’s decision.

B) Discussion of the rejection of claims 1-40 under 35 U.S.C. § 103(a) as being unpatentable over Hughes and Swami.

Non-Analogous Art/Impermissible Hindsight

With regard to claims 1-40, Appellant respectfully submits that it is generally improper to attempt to combine non-analogous references in an attempt to make a showing of obviousness. *In re Oetiker*, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992) (holding “[t]he combination of elements from non-analogous sources, in a manner that reconstructs the applicant’s invention only with the benefit of hindsight, is insufficient to present a *prima facie* case of obviousness”). Hughes and Swami come from vastly disparate technical arts and could not reasonably be combined by one with ordinary skill in the art.

⁸ *Interconnect Planning Corp. v. Feil*, 774 F.2d 1132, 1143, 227 U.S.P.Q. 543, 551 (Fed. Cir. 1985).

⁹ See MPEP § 2141.02.

¹⁰ *In re Mills*, 16 USPQ2d 1430 (Fed. Cir. 1990); M.P.E.P. § 2143.01.

In particular, Hughes relates to “load/store units within processors.”¹¹ Swami relates to network communications. In the Abstract, Swami is described as disclosing:

[a] system, apparatus, and method for determining network capacity and managing network congestion in response to a change in an end-to-end communication path between sender and receiver hosts.¹²

In contrast, the pending claims relate to access of memory by a processor. For example, claim 9 recites “a functional unit to attempt to access data from memory coupled to the processor based on an access request, wherein the functional unit is to retry attempts to access of the data based on other access requests after receipt of a negative acknowledgement in response to the attempt to access the data.” Applicant submits that the art of network communications is nonanalogous to the art of memory access by a processor. MPEP 2141.01(a) sets forth an example of different uses for a particular type of memory that were considered nonanalogous. Specifically, MPEP 2141.01(a) cited *Wang* as an example of nonanalogous art for the electrical arts. In *Wang*, single in-line memory modules (SIMMS) “for installation on a printed circuit motherboard for use in personal computers” was not analogous to SIMMS in an industrial controller. Thus, if two different fields of use for a same type of memory are nonanalogous, Applicant submits that art related to memory access is not analogous to art for network communications. Furthermore, the cited references solve different problems. Hughes solves a problem related to deadlocking in a multi-processor environment. Swami solves a problem related to managing network congestion. Thus, the references are not solving a similar problem.

Because Hughes and Swami are non-analogous art such that impermissible hindsight would be required for their combination, Appellant respectfully submits that no *prima facie* case of obviousness exists with respect to these claims. Accordingly, Appellant respectfully requests reversal of this basis of rejection of claims 1-40.

¹¹ Hughes at column 1, lines 6-7.

¹² Swami at Abstract.

Combining the References Does Not Teach All Limitations:

Claims 1-3, 5-6, 8-10, 12-15, 17-20 and 27-40

With regard to claim 1, among the differences, claim 1 recites “a congestion detection logic to output a signal that indicates that the resource is congested based on receipt of a consecutive number of negative acknowledgments in response to access requests to the resource.” In the Office Action mailed on 3/23/07, the Office indicated the following:

Hughes et al. fail to disclose the congestion detection logic of claim 1. Swami teaches congestion detection logic to output a signal that indicates that the resource is congested based on receipt of a consecutive number of negative acknowledgements in response to access requests to the resource, in pars. 79 and 82.¹³

Applicant respectfully traverses this assertion. This section of Swami does not disclose or suggest congestion based on receipt of a number of negative acknowledgements in response to access requests. In contrast, this section of Swami relates to detection of congestion if a number of duplicate acknowledgements (ACKs) are received (not negative ACKs).

If it is determined 822 that the ACK is a duplicate ACK, it is determined 832 whether a state variable `DUPLICATE_COUNT` has reached the threshold, `DUP_THRESHOLD`.¹⁴

Moreover, this section of Swami does not disclose or suggest this congestion is in response to access requests to a resource. Rather, the system of Swami relates to transmission of ACKs in response to receipt of a packet by a receiver from a sender of a network. The receiver is not transmitting an ACK in response to access request of a resource.

In order to compute such an estimate, the sender may gradually increase the number of packets sent into the network, and may then rely on the

¹³ Office Action at page 3.

¹⁴ Swami at [0082].

acknowledgements (ACKs) received in order to dynamically adjust to the changes in the network.¹⁵

Thus, the cited references do not disclose or suggest all of the claim limitations. Accordingly, Applicant respectfully submits that the rejection of claim 1 under 35 U.S.C. §103 has been overcome. Because claims 2-3 depend from and further define claim 1, Applicant respectfully submits that the rejection of claims 2-3 has been overcome for at least the same reason.

With regard to claim 5, among the differences, claim 5 recites “a congestion detection logic to detect congestion of access of the data based on receipt of a consecutive number of negative acknowledgments that exceed a threshold prior to access of the data.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 5 under 35 U.S.C. §103 has been overcome. Because claim 6 depends from and further defines claim 5, Applicant respectfully submits that the rejection of claim 6 has been overcome for at least the same reason.

With regard to claim 8, among the differences, claim 8 recites “a congestion control logic to disable the functional unit from attempts to access the cache line for a time period after congestion is detected. Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 8 under 35 U.S.C. §103 has been overcome. Because claims 9-10 depends from and further defines claim 8, Applicant respectfully submits that the rejection of claims 9-10 has been overcome for at least the same reason.

With regard to claim 12, among the differences, claim 12 recites “wherein the first processor includes a congestion detection logic to detect congestion of access to the data based

¹⁵ Swami at [0033].

on receipt of a consecutive number of negative acknowledgements in response to the access requests.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations.

Therefore, Applicant respectfully submits that the rejection of claim 12 under 35 U.S.C. §103 has been overcome. Because claims 13-15 depends from and further define claim 12, Applicant respectfully submits that the rejection of claims 13-15 has been overcome for at least the same reason.

With regard to claim 17, among the differences, claim 17 recites “a congestion detection logic to detect congestion of access of the resource based on a consecutive number of negative acknowledgements received in response to the access requests prior to receipt of a positive acknowledgment in response to one of the access requests within a first time period.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 17 under 35 U.S.C. §103 has been overcome. Because claims 18-20 depends from and further define claim 17, Applicant respectfully submits that the rejection of claims 18-20 has been overcome for at least the same reason.

With regard to claim 27, among the differences, claim 27 recites “detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 27 under 35 U.S.C. §103 has been overcome. Because claims 28-30 depends from and further define claim 27, Applicant respectfully submits that the rejection of claims 28-30 has been overcome for at least the same reason.

With regard to claim 31, among the differences, claim 31 recites “detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time

period, prior to receiving a positive acknowledgments.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 31 under 35 U.S.C. §103 has been overcome. Because claims 32-33 depends from and further define claim 31, Applicant respectfully submits that the rejection of claims 32-33 has been overcome for at least the same reason.

With regard to claim 34, among the differences, claim 34 recites “detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 34 under 35 U.S.C. §103 has been overcome. Because claims 35-37 depends from and further define claim 34, Applicant respectfully submits that the rejection of claims 35-37 has been overcome for at least the same reason.

With regard to claim 38, among the differences, claim 38 recites “detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time period, prior to receiving a positive acknowledgments.” Based on the remarks set forth above regarding claim 1, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 38 under 35 U.S.C. §103 has been overcome. Because claims 39-40 depends from and further define claim 38, Applicant respectfully submits that the rejection of claims 39-40 has been overcome for at least the same reason.

Claims 4, 7, 11, 16, 21

In addition to the remarks set forth above regarding claim 1 from which claim 4 depends, Applicant respectfully submits the following remarks. With regard to claim 4, among the differences, claim 4 recites “wherein the congestion control logic is to exponentially decrease the delay after the congestion detection logic receive a number of positive acknowledgements in response to access requests to the resource.” In the Office Action mailed on 3/23/07, the Office indicated that this limitation is disclosed by Swami at [0066]. Applicant respectfully traverses this assertion. This section of Swami relates to exponentially increasing the amount of data prior to receiving of an acknowledgement.:

The cwnd is a state variable that refers to the sender-side limit of the quantity of data that the sender can transmit into the network before receiving an acknowledgement (ACK). . . This is referred to as the "slow start" phase, at which time cwnd may grow exponentially. (emphasis added).¹⁶

Thus, this section of Swami relates to increasing of data – not decrease in delay. Thus, the cited references do not disclose or suggest all of the claim limitations. Accordingly, Applicant respectfully submits that the rejection of claim 4 under 35 U.S.C. §103 has been overcome.

With regard to claim 7, among the differences, claim 7 recites “wherein the congestion control logic is to exponentially decrease the time period after the congestion detection logic receives a number of positive acknowledgements in response to attempts to access data in the memory.” Based on the remarks set forth above regarding claim 4, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 7 under 35 U.S.C. §103 has been overcome.

With regard to claim 11, among the differences, claim 11 recites “wherein the congestion control logic is to exponentially decrease the time period after the congestion detection logic receives a number of positive acknowledgements in response to attempts to access other cache lines in the cache memory.” Based on the remarks set forth above regarding claim 4, Applicant

¹⁶ Swami at [0065] – [0066].

respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 11 under 35 U.S.C. §103 has been overcome.

With regard to claim 16, among the differences, claim 16 recites “wherein the congestion control logic is to disable the first processor from transmitting the access requests for a time period, wherein the time period is based on an exponential back off delay operation.” Based on the remarks set forth above regarding claim 4, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 16 under 35 U.S.C. §103 has been overcome.

With regard to claim 21, among the differences, claim 21 recites “wherein the congestion control logic is to disable the load/store unit from attempts to access the resource for a second time period, wherein the second time period is based on an exponential back off delay operation.” Based on the remarks set forth above regarding claim 4, Applicant respectfully submits that the cited references do not disclose or suggest all of the claim limitations. Therefore, Applicant respectfully submits that the rejection of claim 21 under 35 U.S.C. §103 has been overcome.

Claims 22-26

With regard to claim 22, among the differences, claim 22 recites “a congestion detection logic to detect congestion of access of a first cache line of the number of cache lines based on a ratio of a number of negative acknowledgments to a number of positive acknowledgments received in response to the access requests.” In addition to the remarks regarding claim 1, Applicant submits the following remarks.

In the Office Action mailed on 3/23/07, the Office indicated that “Swami teaches that the congestion logic can be any statistical parameter using negative and positive acknowledgements, in pars. 79-82.”¹⁷ Applicant respectfully traverses this assertion. As noted above, this section of Swami is limited to detection based on a number of duplicate ACKs. This section of Swami

¹⁷ Office Action at page 11.

does not disclose or suggest the use of negative ACKs for detection of congestion. Moreover, this section of Swami does not disclose or suggest any type of statistical parameter beyond the tracking of duplicate ACKs.

Thus, the cited references do not disclose or suggest all of the claim limitations. Accordingly, Applicant respectfully submits that the rejection of claim 22 under 35 U.S.C. §103 has been overcome. Because claims 23-26 depend from and further define claim 22, Applicant respectfully submits that the rejection of claims 23-26 has been overcome for at least the same reason.

8. SUMMARY

It is respectfully submitted that the claims are patentable over the cited art. Reversal of the rejection and allowance of the pending claim are respectfully requested.

Respectfully submitted,

GREGORY MARLAN et al.

By their Representatives,

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

P.O. Box 2938

Minneapolis, MN 55402

Date

8-9-07

By



Gregg A. Peacock

Reg. No. 45,001

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: Mail Stop Amendment, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 9 day of August 2007.



Name



Signature

CLAIMS APPENDIX

1. (Original) An apparatus comprising:

a load/store unit that includes a retry logic that is to retry access to a resource after receipt of a negative acknowledgement for an attempt to access the resource by the load/store unit; and

a congestion detection logic to output a signal that indicates that the resource is congested based on receipt of a consecutive number of negative acknowledgments in response to access requests to the resource.

2. (Original) The apparatus of claim 1 further comprising a congestion control logic to disable the retry logic from retry accesses to the resource based on receipt of the signal from the congestion detection logic.

3. (Original) The apparatus of claim 2, wherein the congestion control logic is to exponentially increase the delay after the congestion detection logic is to detect congestion while the resource is currently congested.

4. (Original) The apparatus of claim 2, wherein the congestion control logic is to exponentially decrease the delay after the congestion detection logic receive a number of positive acknowledgements in response to access requests to the resource.

5. (Original) A processor comprising:

a functional unit to attempt to access data from memory coupled to the processor based on an access request, wherein the functional unit is to retry attempts to access of the data based on other access requests after receipt of a negative acknowledgement in response to the attempt to access the data; and

a congestion detection logic to detect congestion of access of the data based on receipt of a consecutive number of negative acknowledgments that exceed a threshold prior to access of the data; and

a congestion control logic to disable the functional unit from the attempts to access the data for a time period after congestion is detected.

6. (Original) The processor of claim 5, wherein the congestion control logic is to exponentially increase the time period after the congestion detection logic is to detect congestion while access to other data in the memory is congested.

7. (Original) The processor of claim 6, wherein the congestion control logic is to exponentially decrease the time period after the congestion detection logic receives a number of positive acknowledgements in response to attempts to access data in the memory.

8. (Original) A processor comprising:

a functional unit to attempt to access a cache line in a cache memory coupled to the processor based on an access request, wherein the functional unit is to retry attempts to access the cache line based on additional access requests after receipt of a negative acknowledgement in response to the attempt to access the data;

a congestion detection logic to detect congestion of access of the cache line based on an average number of negative acknowledgments received that exceed a threshold prior to access of the data; and

a congestion control logic to disable the functional unit from attempts to access the cache line for a time period after congestion is detected.

9. (Original) The processor of claim 8, wherein the average number of negative acknowledgements is within a window and wherein the congestion detection logic is to move the window over time of attempts to access the cache line by the functional unit.

10. (Original) The processor of claim 8, wherein the congestion control logic is to exponentially increase the time period after the congestion detection logic is to detect congestion while access of other cache lines in the cache memory is congested.

11. (Original) The processor of claim 8, wherein the congestion control logic is to exponentially decrease the time period after the congestion detection logic receives a number of positive acknowledgements in response to attempts to access other cache lines in the cache memory.

12. (Original) A system comprising:
a cache memory to store data; and
a first processor to attempt to access the data from the cache memory based on access requests, wherein the first processor includes a congestion detection logic to detect congestion of access to the data based on receipt of a consecutive number of negative acknowledgements in response to the access requests.

13. (Original) The system of claim 12 further comprising:
a second processor associated with the cache memory;
a hub controller to receive the access requests from the first processor, the hub controller to forward the access requests to the second processor, wherein the second processor is to determine whether the data in the cache memory is accessible.

14. (Original) The system of claim 13, wherein the second processor is to transmit a negative acknowledgement back to the first processor through the hub controller if the data is not accessible, the second processor to transmit a positive acknowledgment back to the first processor through the hub controller if the data is accessible.

15. (Original) The system of claim 12, wherein the first processor further comprises a congestion control logic to disable the first processor from transmitting the access requests if the congestion detection logic determines that access to the data is congested.

16. (Original) The system of claim 12, wherein the congestion control logic is to disable the first processor from transmitting the access requests for a time period, wherein the time period is based on an exponential back off delay operation.

17. (Original) A system comprising:

a resource; and

a first processor having a load/store functional unit, the load/store functional unit to attempt to access the resource based on access requests, wherein the first processor includes a congestion detection logic to detect congestion of access of the resource based on a consecutive number of negative acknowledgements received in response to the access requests prior to receipt of a positive acknowledgment in response to one of the access requests within a first time period.

18. (Original) The system of claim 17 further comprising:

a second processor associated with the resource;

a hub controller to receive the access requests from the first processor, the hub controller to forward the access requests to the second processor, wherein the second processor is to determine whether the resource is accessible.

19. (Original) The system of claim 18, wherein the second processor is to transmit a negative acknowledgement back to the first processor through the hub controller if the resource is not accessible, the second processor to transmit a positive acknowledgment back to the first processor through the hub controller if the resource is accessible.

20. (Original) The system of claim 17, wherein the first processor further comprises a congestion control logic to disable the load/store functional unit from attempting to access the resource if the congestion detection logic is to detect congestion of access of the resource

21. (Original) The system of claim 17, wherein the congestion control logic is to disable the load/store unit from attempts to access the resource for a second time period, wherein the second time period is based on an exponential back off delay operation.

22. (Original) A system comprising:

a cache memory to include a number of cache lines for storage of data; and
at least two processors, wherein a first processor of the at least two processors is to attempt to access the data in one of the number of cache lines based on access requests, wherein the first processor includes a congestion detection logic to detect congestion of access of a first cache line of the number of cache lines based on a ratio of a number of negative acknowledgments to a number of positive acknowledgments received in response to the access requests.

23. (Original) The system of claim 22, wherein a second processor of the at least two processors is associated with the cache memory and wherein the system further comprises a hub controller, the hub controller to receive the access requests from the first processor, the hub controller to forward the access requests to the second processor, wherein the second processor is to determine whether the one of the number of cache lines is accessible.

24. (Original) The system of claim 23, wherein the second processor is to transmit a negative acknowledgement back to the first processor through the hub controller if the one of the number of cache lines is not accessible, the second processor to transmit a positive acknowledgement back to the first processor through the hub controller if the one of the number of cache lines is accessible.

25. (Original) The system of claim 22, wherein the first processor further comprises a congestion control logic to disable, for a time period, the first processor to attempt to access the data if the congestion detection logic is to detect congestion of access of the first cache line.

26. (Original) The system of claim 25, wherein the congestion control logic is to exponentially increase the time period after the congestion detection logic is to detect congestion while access to other cache lines in the cache memory.

27. (Previously Presented) A method comprising:
transmitting access requests, by a first processor, to access data in a memory;
receiving, by the first processor, a positive acknowledgement or a negative acknowledgment from a second processor that is associated with the memory based on one of the number of access requests; and

detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment.

28. (Original) The method of claim 27 further comprising controlling access to the data in the memory if the consecutive number of negative acknowledgements, received by the first processor, exceeds the first threshold, prior to receipt of the positive acknowledgment.

29. (Original) The method of claim 28, wherein controlling access to the data in the memory comprises disabling transmitting of the access requests, by the first processor, for a time period.

30. (Original) The method of claim 29, wherein controlling access to the resource comprises exponentially increasing the time period upon determining that the congestion is detected for other data in the memory while the time period has not expired.

31. (Original) A method comprising:
accessing, by at least one processor, a resource based on an access request;
receiving a positive acknowledgement if the resource is accessible;
receiving a negative acknowledgement if the resource is not accessible;
retrying accessing, by the at least one processor, of the resource based on a number of access requests; and

detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time period, prior to receiving a positive acknowledgments.

32. (Original) The method of claim 31 further comprising controlling access to the resource if the consecutive number of negative acknowledgements, received by the at least one processor, exceeds the first threshold, prior to receipt of the positive acknowledgement.

33. (Original) The method of claim 31, wherein controlling access to the resource comprises disabling transmitting of the access requests, by the first processor, for a time period.

34. (Previously Presented) A computer storage medium that provides instructions, which when executed by a machine, cause said machine to perform operations comprising:
transmitting access requests, by a first processor, to access data in a memory;
receiving, by the first processor, a positive acknowledgement or a negative acknowledgment from a second processor that is associated with the memory based on one of the number of access requests; and

detecting congestion of the data based on receipt, by the first processor, of a consecutive number of negative acknowledgements that exceed a first threshold, prior to receipt, by the first processor, of a positive acknowledgment.

35. (Previously Presented) The computer storage medium of claim 34 further comprising controlling access to the data in the memory if the consecutive number of negative acknowledgements, received by the first processor, exceeds the first threshold, prior to receipt of the positive acknowledgement.

36. (Previously Presented) The computer storage medium of claim 35, wherein controlling access to the data in the memory comprises disabling transmitting of the access requests, by the first processor, for a time period.

37. (Previously Presented) The computer storage medium of claim 36, wherein controlling access to the resource comprises exponentially increasing the time period upon determining that the congestion is detected for other data in the memory while the time period has not expired.

38. (Previously Presented) A computer storage medium that provides instructions, which when executed by a machine, cause said machine to perform operations comprising:

- accessing, by at least one processor, a resource based on an access request;
- receiving a positive acknowledgement if the resource is accessible;
- receiving a negative acknowledgement if the resource is not accessible;
- retrying accessing, by the at least one processor, of the resource based on a number of access requests; and

- detecting that a consecutive number of negative acknowledgements exceeds a first threshold within a time period, prior to receiving a positive acknowledgments.

39. (Previously Presented) The computer storage medium of claim 38 further comprising controlling access to the resource if the consecutive number of negative acknowledgements, received by the at least one processor, exceeds the first threshold, prior to receipt of the positive acknowledgement.

40. (Previously Presented) The computer storage medium of claim 39, wherein controlling access to the resource comprises disabling transmitting of the access requests, by the first processor, for a time period.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.